IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application for:

BATCH MODE SESSION-BASED ENCRYPTION OF VIDEO ON DEMAND CONTENT

Inventor(s): Leo Mark Pedlow, Jr. and Davender Agnihotri

Docket Number: SNY-T5709.02

Prepared By:

Miller Patent Services
2500 Dockery Lane
Raleigh, NC 27606

Phone: (919) 816-9981
Fax: (919) 816-9982
Email: miller@patent-inventions.com

CERTIFICATE OF EXPRESS MAILING FOR NEW PATENT APPLICATION

"Express Mail" mailing label number: ER999163752US

Date of Deposit: 4/21/2004

I Hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner for Patents, PO Box 1450, Alexandria, VA 22313-1450.

Typed or printed name of person mailing paper or fee: Catherine N. Miller

Signature of person mailing paper or fee: Catherine N. Miller

BATCH MODE SESSION-BASED ENCRYPTION OF VIDEO ON DEMAND CONTENT

CROSS REFERENCE TO RELATED DOCUMENTS

This application is related to and claims priority benefit of U.S. Provisional Patent Application Serial No. 60/516,131 filed October 31, 2003 to Pedlow et al. for "Batch Mode Session Based Encryption of Video On Demand Content" which is hereby incorporated by reference; this application also claims priority benefit of the following pending U.S. Patent Applications:

Serial Number 10/764,202, filed January 23, 2004, Case Number SNY-T5708.01;

Serial Number 10/764,011, filed January 23, 2004, Case Number SNY-T5710.01;

Serial Number 10/802,084, filed March 16, 2004, Case Number SNY-T5711.02;

Serial Number 10/802,007, filed March 16, 2004, Case Number SNY-T5712.02;

Serial Number 10/802,008, filed March 16, 2004, Case Number SNY-T5717.02; and

Serial Number 10/823,431, filed April 13, 2004, Case Number SNY-T5775.02, each of which names Leo Mark Pedlow, Jr. as an inventor and which are hereby incorporated by reference herein.

This application is also related to U.S. Patent Applications docket number SNY-R4646.01 entitled "Critical Packet Partial Encryption" to Unger et al., serial number 10/038,217; patent applications docket number SNY-R4646.02 entitled "Time Division Partial Encryption" to Candelore et al., serial number 10/038,032; docket number SNY-R4646.03 entitled "Elementary Stream Partial Encryption" to Candelore, serial number 10/037,914; docket number SNY-R4646.04 entitled "Partial Encryption and PID Mapping" to Unger et al., serial number 10/037,499; and docket number SNY-R4646.05 entitled "Decoding and Decrypting of Partially Encrypted Information" to Unger et al.,

serial number 10/037,498 all of which were filed on January 2, 2002 and are hereby incorporated by reference herein.

COPYRIGHT NOTICE

BACKGROUND

The Passage™ initiative (Passage is a trademark of Sony Electronics Inc.), promoted by Sony, provides a mechanism for MSOs (Multiple Service Operators) to deploy non-legacy headend equipment, subscriber devices and services on their existing legacy networks. In the USA, these networks are most commonly supplied by either Motorola (formerly General Instrument) or Scientific Atlanta. These two companies at present constitute better than a 99% share of the US cable system market as turnkey system providers. The systems, by design, employ proprietary technology and interfaces precluding the introduction of non-incumbent equipment into the network. An MSO, once choosing one of these suppliers during conversion from an analog cable system to a digital cable system, faces a virtual monopoly when seeking suppliers for additional equipment as their subscriber base or service offering grows.

Before the Passage™ initiative, the only exit from this situation was to forfeit the considerable capital investment already made with the incumbent provider, due to the intentional incompatibility of equipment between the incumbent and other sources. One primary barrier to interoperability is in the area of conditional access systems, the heart of addressable subscriber management and revenue collection resources in a modern digital cable network.

The Passage™ technologies were developed to allow the independent coexistence of two or more conditional access systems on a single, common plant. Unlike other

attempts to address the issue, the two systems operate with a common transport stream without any direct or indirect interaction between the conditional access systems. The basic processes used in these technologies are discussed in detail in the above-referenced pending patent applications.

The above-referenced commonly owned patent applications, and others, describe inventions relating to various aspects of methods generally referred to herein as partial encryption or selective encryption, consistent with certain aspects of Passage™. More particularly, systems are described therein wherein selected portions of a particular selection of digital content are encrypted using two (or more) encryption techniques while other portions of the content are left unencrypted. By properly selecting the portions to be encrypted, the content can effectively be encrypted for use under multiple decryption systems without the necessity of encryption of the entire selection of content. In some embodiments, only a few percent of data overhead is consumed to effectively encrypt the content using multiple encryption systems. This results in a cable or satellite system being able to utilize Set-top boxes (STB) or other implementations of conditional access (CA) receivers from multiple manufacturers in a single system - thus freeing the cable or satellite company to competitively shop for providers of Set-top boxes.

In each of these disclosures, the clear content is identified using a primary Packet Identifier (PID). A secondary PID (or shadow PID) is also assigned to the program content. Selected portions of the content are encrypted under two (or more) encryption systems and the encrypted content transmitted using both the primary and secondary PIDs (one PID or set of PIDs for each encryption system). The so-called legacy STBs operate in a normal manner decrypting encrypted packets arriving under the primary PID and ignoring secondary PIDs. The newer (non-legacy) STBs operate by associating both the primary and secondary PIDs with a single program. Packets with a primary PID are decoded normally and packets with a secondary PID are first decrypted then decoded. The packets associated with both PIDs are then assembled together to make up a single program stream. The PID values associated with the packets are generally remapped to a single PID value for decoding (shadow PIDs remapped to the primary PID value or vice versa.)

BRIEF DESCRIPTION OF THE DRAWINGS

Certain illustrative embodiments illustrating organization and method of operation, together with objects and advantages may be best understood by reference detailed description that follows taken in conjunction with the accompanying drawings in which:

FIGURE 1 is a block diagram of a clear video VOD system.

FIGURE 2 is a diagram illustrating storage of I-frame data to support trick mode operation in a VOD system.

FIGURE 3 is a block diagram of a pre-encrypted VOD system using a single (legacy) encryption system.

FIGURE 4 is a block diagram depicting a hybrid composite VOD system architecture consistent with certain embodiments of the present invention.

FIGURE 5 is a block diagram of a re-encrypted VOD architecture consistent with certain embodiments of the present invention.

FIGURE 6 illustrates a dynamic composition pre-encrypted VOD architecture consistent with certain embodiments of the present invention.

FIGURE 7 illustrates a dynamic composition pre-encrypted VOD architecture using dual trick play indices consistent with certain embodiments of the present invention.

FIGURE 8 is a block diagram of a segregated session based encrypted VOD architecture consistent with certain embodiments of the present invention.

FIGURE 9 is a block diagram of a composite session based encrypted VOD architecture consistent with certain embodiments of the present invention.

FIGURE 10 illustrates composite session based encryption content flow consistent with certain embodiments of the present invention.

FIGURE 11 illustrates batch based encrypted VOD server content flow consistent with certain embodiments of the present invention.

FIGURE 12 illustrates an optimized batch based encrypted VOD server content flow consistent with certain embodiments of the present invention.

FIGURE 13 is a flow chart depicting batch based encrypted VOD.

ACRONYMS, ABBREVIATIONS AND DEFINITIONS

ASI - Asynchronous Serial Interface

CA - Conditional Access

CASID - Conditional Access System Identifier

CPE - Customer Premises Equipment

DHEI - Digital Headend Extended Interface

ECM - Entitlement Control Message

EPG - Electronic Program Guide

GOP - Group of Pictures (MPEG)

MPEG - Moving Pictures Experts Group

MSO - Multiple System Operator

OLES – Off Line Encryption System

PAT - Program Allocation Table

PID - Packet Identifier

PMT - Program Map Table

PSI - Program Specific Information

QAM - Quadrature Amplitude Modulation

RAID – Redundant Array of Independent Disks

RAM - Random Access Memory

SAN - Storage Area Network

VOD - Video on Demand

Critical Packet - A packet or group of packets that, when encrypted, renders a portion of a video image difficult or impossible to view if not properly decrypted, or which renders a portion of audio difficult or impossible to hear if not properly decrypted. The term "critical" should not be interpreted as an absolute term, in that it may be possible to hack an elementary stream to overcome encryption of a "critical packet", but when subjected to normal decoding, the inability to fully or properly decode such a "critical packet" would inhibit normal viewing or listening of the program content.

Selective Encryption (or Partial Encryption) – encryption of only a portion of an elementary stream in order to render the stream difficult or impossible to use (i.e., view or hear).

Dual Selective Encryption – encryption of portions of a single selection of content under two separate encryption systems.

Passage™ - Trademark of Sony Electronics Inc. for various single and multiple selective encryption systems, devices and processes.

Trick mode – an operational mode of playback of digital content to simulate fast forward, rewind, pause, suspend (stop), slow motion, etc. operations as in a video tape system.

The terms "a" or "an", as used herein, are defined as one, or more than one. The term "plurality", as used herein, is defined as two or more than two. The term "another", as used herein, is defined as at least a second or more. The terms "including" and/or "having", as used herein, are defined as comprising (i.e., open language). The term "coupled", as used herein, is defined as connected, although not necessarily directly, and not necessarily mechanically. The term "program", as used herein, is defined as a sequence of instructions designed for execution on a computer system. A "program", or "computer program", may include a subroutine, a function, a procedure, an object method, an object implementation, in an executable application, an applet, a servlet, a source code, an object code, a shared library / dynamic load library and/or other sequence of instructions designed for execution on a computer system.

The terms "scramble" and "encrypt" and variations thereof may be used synonymously herein. Also, the term "television program" and similar terms can be interpreted in the normal conversational sense, as well as a meaning wherein the term means any segment of A/V content that can be displayed on a television set or similar monitor device. The term "video" is often used herein to embrace not only true visual information, but also in the conversational sense (e.g., "video tape recorder") to embrace not only video signals but associated audio and data. The term "legacy" as used herein refers to existing technology used for existing cable and satellite systems. The exemplary embodiments of VOD disclosed herein can be decoded by a television Set-Top Box

(STB), but it is contemplated that such technology will soon be incorporated within television receivers of all types whether housed in a separate enclosure alone or in conjunction with recording and/or playback equipment or Conditional Access (CA) decryption module or within a television set itself.

DETAILED DESCRIPTION

While this invention is susceptible of embodiment in many different forms, there is shown in the drawings and will herein be described in detail specific embodiments, with the understanding that the present disclosure of such embodiments is to be considered as an example of the principles and not intended to limit the invention to the specific embodiments shown and described. In the description below, like reference numerals are used to describe the same, similar or corresponding parts in the several views of the drawings.

CLEAR VOD ARCHITECTURES

The decision on a particular VOD architecture is the result of the interaction between a complex set of both independent and dependent variables, providing a solution to an equation of state. Some of the variables are fixed directly as a result of choices by the MSO. Others are constrained by factors such as the existing incumbent system, location, size, available capital and return on investment requirements.

A generalized VOD system 10, as shown in **FIGURE 1**, contains some or all of the following elements / resources: Content Aggregation and Asset management 14, Content distribution (SAN) 18, Video server module(s) 22, Session Management 26, Transaction management 30, Billing system 34, EPG server or VOD catalog server 38, Transport router/switch fabric (routing matrix) 42, Stream encryption device(s) (not shown in this Figure), and QAM modulators/upconverters and other edge resources 46. This VOD system 10 provides programming to the subscriber terminals such as 50 for ultimate viewing and listening on a TV set or other monitor device 54.

In operation, content is received from various sources including, but not limited to, satellite broadcasts received via one or more satellite dishes 58. Content is aggregated

at 14 and cataloged at EPG server or VOD catalog server 38. Content is then distributed at 18 to one or more video servers 22. When a subscriber orders a VOD selection, a message is sent from the subscriber terminal (e.g., STB) 50 to the session manager 26. The session manager 26 notifies the transaction manager 30 to assure that the billing system 34 is properly brought into play. The session manager 26 selects a VOD server from a cluster of VOD servers having the requested content on it and having a signal path that reaches the node serving the subscriber. The session manager also enables the routing matrix 42 to properly route the selected video content through the correct edge resources 46 for delivery to the subscriber terminal 50.

TRICK MODES

One aspect of VOD that has become a "signature" feature is the support of "trick modes". These are operational modes invoked by the session client that mimic a traditional VCR or DVD player and includes fast forward, rewind, pause, suspend (stop), slow motion, etc. Trick modes have been heretofore implemented through the creation of multiple files containing a subset of the original content (subfiles) as illustrated in **FIGURE 2**. The content is generally stored in a set of RAID drives 70. A particular selection of content is stored in its entirety in a file 74 within the RAID drives 70. A set of subfiles for rewind and fast forward trick modes (files 78 and 80 respectively) contain I-frames ordered in a manner that will permit playback sequentially to achieve the rewind and fast forward effect. Generally, these subfiles contain only I-frames, since I-frames contain stand-alone whole pictures (see ISO/IEC 13818-2, section 6.1.1.7). I-frames are somewhat larger than B or P frames, and they generally represent approximately 21% of the data in a given video selection.

A file containing only I-frames extracted from the original content affords the ability to have accelerated playback, since typical GOP (group of pictures) structures have only one frame in about 10 to 20 as an I-frame. If the I-frame files are played at normal rates (1 frame per 33 mS) the pictures will appear to the viewer to sequence at about a 10x to 20x rate, though the actual data rate is the same as the original content. If

the I-frame sequence is reversed in the file, the motion will appear to run backwards. This is the method used to implement fast forward and rewind.

By attaching an index count to match the I-frames in the original content file to the duplicated I-frames stored in the associated subfiles 78 and 80, a method is provided to allow immediate transition from normal speed forward play to fast forward or rewind. In operation the video server plays the selected content file and upon subscriber selection of a trick mode (or vice versa) the server notes the index value of the closest I-frame and then opens the appropriate associated subfile 78 or 80 and moves to the I-frame in the subfile with the same corresponding index. The video server treats all stream content (main file or subfiles) the same and always spools the MPEG packets to the outgoing transport stream at the same constant bit rate through multiplexers and buffers 84 as shown. It is through this method that trick modes are typically implemented on a slotted, session based system without the encumbrance of additional, dynamic bit rate issues.

Unfortunately, the use of such multiple subfiles results in storage space inefficiencies. As will be seen, these inefficiencies can become compounded in systems utilizing multiple encryption.

VOD PROGRAM SPECIFIC INFORMATION

A function of the VOD video server(s) 22, in addition to origination of session A/V content, is the creation of the associated, session specific PSI (program specific information). This information is a departure from the broadcast model in that the PSI is extremely dynamic. The content of the PAT and subordinate PMTs change whenever a new session is started or ended. In the broadcast world, the PSI changes very seldom because the PSI tables reflect only the structure of the transport multiplex, not the actual A/V content carried within.

The VOD video server 22 dynamically assigns a new session to an existing, available "slot" in an outgoing transport multiplexed stream. The slot is denoted by the MPEG program number and in many cases, the combination of which transport stream (TSID) and program number determine at the service level a unique session and the routing that occurs as a result. Edge resources 46 generally are not configured

dynamically. The routing of content appearing on a particular input port to a specific QAM carrier at the output is determined through a preconfigured, static assignment of TSID/input port and program number mapping to specific QAM resources in the device. This same mapping information is also loaded in the VOD system so that once a session is requested by and authorized for a specific subscriber terminal 50, a solution to a routing matrix 42 can be determined to find the appropriate VOD server 22 and QAM transport 46 serving the requestor. This solution also considers dynamic issues such as which servers 22 the requested asset is loaded upon, and server loading/available slots in addition to the simpler, static solution to finding the first possible path to the requesting subscriber terminal 50.

In addition to solving the routing matrix 42 and provisioning the session with PIDs and PSI appropriate to follow the intended route, elements of the same information (program ID and QAM frequency) are also communicated to the session client at subscriber terminal 50 at the subscriber's premises so that the requested stream can be properly received and presented to the subscriber.

CLEAR VOD DISTRIBUTION

Perhaps the simplest VOD implementation is a clear VOD distribution system, i.e. one that contains no encryption as depicted in **FIGURE 1**. While not providing any safekeeping of what might be considered the entertainment medium's most valuable properties, namely current feature films, etc., clear VOD avoids many of the issues that the incumbent cable system providers to date have not adequately addressed and that introduction of a second, alternative CA system complicates even further still. Various arrangements for providing selective or full encryption in a VOD environment are discussed below. Throughout this discussion, it is instructive to carry an example VOD movie through the various embodiments to illustrate the relative storage efficiencies obtained with the various systems disclosed. A real world example of a VOD movie which will be used throughout this document has the following attributes:

Compressed video data rate: 3Mbit/S

Movie length: 120 minutes (2 Hrs)

I-frame overhead:	17%
Total storage used for the video portion of a single, clear (unencrypted) copy of a film:	**3.618GBytes.**

PRE-ENCRYPTED VOD DISTRIBUTION

Pre-encrypted VOD systems such as system 100 shown in **FIGURE 3** can be architecturally similar to clear VOD distribution systems. One difference between the two is that on pre-encrypted systems there is pre-processing of the content prior to storage in the VOD system to provide safekeeping of content during the storage and distribution phases. This pre-processing is carried out in pre-encryptor 104. Data security is implemented through storage of previously encrypted content within the video server(s) 22. While the clear VOD system contains directly viewable MPEG or other compressed A/V content on the server(s) 22, the pre-encrypted model stores this same content in a form that is only decipherable using a properly entitled subscriber terminal 50.

The pre-encryption process can be performed by the MSO at the time of deployment on the VOD system 100, prior to loading into the storage area network (SAN) used to propagate content to all of the video servers in the MSO's system. Alternatively, the encryption may be performed prior to receipt of the content by the MSO at an external service bureau, content aggregator or by the distributor or studio. In this case, the content is theoretically secured throughout the distribution phase, storage phase and transmission to subscriber for display on an authorized device. The use of pre-encryption prior to distribution of content to the MSO potentially adds to the complexity of entitlement distribution, separate from the content distribution, for installation on the VOD transaction manager 30 to allow bone fide subscribers to decrypt the purchased content.

Many pre-encrypted VOD architectures share one or more of the following common drawbacks:

- Additional handling of new content may be needed to perform the pre-encryption prior to loading into the server, either by the MSO or service bureau.
- Coordination and/or distribution is required for entitlements matching the access criteria used to encrypt the content stored in the server.
- Limited "shelf life" of the encryption keys used to secure the stored content, rendering decryption impossible at a later date.
- Incapability of present VOD video servers to load pre-encrypted streams.
- Incompatibility of pre-encrypted streams with present methods supporting trick mode play (fast-forward & rewind) on screen.
- One common key is used for all sessions accessing a particular program and it remains the same for the duration of time the content is in inventory on the server.
- According to MSOs familiar with the subject, pre-encrypted VOD streams are unsupported by conditional access technologies from certain manufacturer(s).

The issue regarding trick play and pre-encryption is based upon the concept that VOD servers 22 currently expect clear content and then subsequently identify the I-frames and store or otherwise segregate them for access in fast-forward or fast rewind playback modes, as described in conjunction with **FIGURE 2**. If the stream is pre-encrypted prior to storage upon the server, it may be difficult or impossible for the server 22 to examine packet payloads to identify I-frames during the process of importation into the server 22 to create trick mode files 78 and 80 or associated indices. Many current systems will not accept streams for importation that are pre-encrypted.

SEGREGATED STORAGE PRE-ENCRYPTION

A segregated storage mechanism can be physically similar to the architecture of the clear VOD distribution system. The content is encrypted in its entirety (100%) and a separate copy of the complete feature is stored for each different conditional access format supported by the MSO. The organization and configuration of the system is such that when a subscriber initiates a session on the server, the stream files for the selected content containing the CA format appropriate to the specific equipment deployed at the

subscriber's premises requesting the session are spooled and delivered. This method offers a low system complexity encrypted VOD system but may suffer from the same issues common to other pre-encryption topologies, mentioned previously. In addition, a very significant storage penalty (one or more encrypted duplicate copies of the same movie) is incurred.

If one refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require an additional 7.236GBytes to store using segregated pre-encryption supporting two different CA systems.

Changes to the method employed by the VOD system are used for creating dynamic PSI data to implement this architecture supporting multiple CA systems. The VOD system session manager is made aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information is in turn transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server is cognizant of the conditional access resources (ECMs) for each program stored on the server and these resources can be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

COMPOSITE STORAGE PRE-ENCRYPTION

Composite storage is essentially the storage on the video server of a selectively encrypted stream such as a Passage™ processed stream that contains previously encrypted "critical packets" for a plurality (two or more) of independent conditional access systems (i.e., dual selective encrypted). The stream may be prepared identically to the processing of a selectively encrypted broadcast stream as described in the above-referenced pending patent applications, except that the resultant transport stream is recorded to a hard disk or other suitable computer readable storage medium, instead of

being sent directly to a QAM modulator for HFC distribution to the requesting subscriber. As with other pre-encryption models, the content can be encrypted by either the MSO at time of deployment on the VOD system, a third party service bureau, by the studios themselves (the latter two cases being prior to receipt of the content by the MSO), or by or under control of other entities.

In this embodiment the small additional overhead in content storage (usually about 2% – 10% representing "critical packets" that are multiple encrypted) is traded for the support of multiple independent CA formats without replication of entire streams. A negative aspect, in addition to those mentioned previously and common to other pre-encryption topologies, is the vulnerability of the prepared selectively encrypted stream to corruption by downstream equipment containing transport re-multiplexing functionality that is not specifically designed to maintain the integrity of the selective encryption process applied to the stream.

If one refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require approximately 3.690GBytes to store using composite storage pre-encryption supporting two different CA systems with a critical packet "density" of 2%.

Certain changes to the method employed by the VOD system for creating dynamic PSI data can be used to implement this architecture. The VOD system session manager can be made to be aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information is in turn transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server is cognizant of the conditional access resources (ECMs) for each program stored on the server and these can be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, can indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

Likewise, the video server dynamically allocates another set of PIDs for the shadow packets associated with the respective audio and video component streams for each session in the manner described in the above-referenced patent applications. This information can be included in the PSI sent in sessions requested by non-legacy clients. In total, eight different PIDs and corresponding data resources are dynamically allocated and managed by the server for each session: PAT (one table common to all sessions, but modified for each), PMT, Primary Video, Primary Audio, Shadow Video, Shadow Audio, Legacy ECM and Alternative ECM. Six of these entities can be stored in the embedded stream and use dynamic PID remapping for each session.

Consider the issue of which device to use in conjunction with performing the legacy encryption of the "critical" packets prior to storage on the VOD video server. If the legacy device is specially designed to process content destined for loading into a VOD video server, it may not accept a selectively encrypted stream at its input. The content format specified for VOD servers often uses a single program transport multiplex containing a single PAT entry, single PMT entry and service components, for one audio and one video stream. The shadow packets added in a composite selectively encrypted transport stream may prove problematic for a legacy VOD pre-encryption device, in certain instances. It is more probable that a device or process (since there are no real time requirements, an off-line process running on a PC or UNIX server may suffice) to process a candidate stream before passing through the legacy pre-encryptor and then post-encryption reconcile to extract only the encrypted "critical" packets for insertion into the VOD video server 22. The same or similar algorithms and techniques for performing this manipulation for selective encryption processing as described in the above-referenced patent applications can be adapted to VOD applications for off-line work.

The VOD server 22 may also be modified to allow introduction of streams having multiple service elements (primary video, primary audio, shadow video, shadow audio) uniquely associated with a Passage™ transport. The present video servers generally only allow one each, primary video and audio, respectively. The quartet of data representing

Passage™ processed A/V content should preferably be managed as a indivisible set on the VOD video server 22.

Some additional bandwidth efficiencies may be obtained if, at the edge resources, shadow packets are removed from the composite streams in sessions serving legacy clients. Similarly, in certain embodiments, the edge resources, if selective encryption aware, could reinsert the shadow packets embedded in the stored stream in place of the legacy encrypted packets on the original program PID. These improvements would result in no carriage overhead for support of multiple conditional access systems on a single transport.

HYBRID COMPOSITE STORAGE PRE-ENCRYPTION

Hybrid composite storage is a variant of the composite storage concept, but incorporates elements of session-based encryption for implementing the alternative conditional access encryption. In this scenario, depicted as system 130 of **FIGURE 4**, the legacy "critical" packets, comprising approximately 2-10% of the total content, are pre-encrypted by the legacy conditional access system 104 using selective encryption technology for managing the process. The selective encryption is managed in selective encryption processor 134. The duplicate copy of "critical" packets, which are located on previously unused PIDs, is left unencrypted. This latter aspect is the departure from the composite storage scenario described above. The composite stream of unencrypted non-critical packets, legacy encrypted "critical" packets on the original service PIDs and an unencrypted, duplicate copy of the "critical" packets on alternate service PIDs is stored on the video server 22 as a single stream.

Upon playback to a subscriber session, if the session is destined for a legacy STB (represented by subscriber terminal 50), the existing paradigm for pre-encrypted content is followed and no special action is taken. The stream is routed at routing matrix 138 operating under control of session manager 26, through a session encryption device 142 capable of performing encryption using the alternative conditional access system 144, but the session manager 26 does not provision the device to perform encryption on elements of the stream and it is sent directly to the requesting subscriber without further

modification. To maintain security of the outgoing stream and to reduce the bandwidth of the session for legacy sessions, the stream is processed through an add-drop remultiplexer 148 and the clear "critical" content on alternate service PIDs are removed from the outgoing transport. The output stream is then routed at routing matrix 152 to appropriate edge resources 46 for delivery to the subscriber terminal 50. In one embodiment, the session encryption device 142 that performs encryption using the alternative conditional access system also contains the add-drop multiplexer capability. Other variations will also occur to those skilled in the art upon consideration of the present teaching.

If, on the other hand, the session is destined for a non-legacy STB (also as represented in this illustration by subscriber terminal 50, the stream is routed through session encryption device 142 capable of performing encryption using the alternative conditional access system and only the "critical" packets on alternate service PIDs (previously in the clear) are encrypted using the alternative conditional access system 144, as provisioned by the session manager.

Some additional bandwidth efficiencies may be obtained for these non-legacy sessions, if the edge device is selective encryption aware, by reinserting the shadow packets embedded in the stored stream, now encrypted, in place of the legacy encrypted packets on the original program PID. This improvement would result in no carriage overhead for support of multiple conditional access systems on a single transport.

A preprocessor can be used to perform selective encryption of content to be loaded onto the video server. A modified file protocol can be used to allow the video server to import and associate these files. Either the preprocessor or the video server can be designed to perform the indexing. An alternate instantiation could be use to perform all selective encryption pre-processing (e.g., PID mapping and packet duplication) within the VOD video server 22 itself. This could be accomplished by modifying the VOD video server 22 application to add a pre-processor task as a separate executable, called by the VOD video server 22 during the process to prepare content for pre-encryption.

Changes can be implemented to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session

manager 26 is made aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information can in turn be transferred to the VOD video server 22 that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The VOD video server 22 is cognizant of the conditional access resources (ECMs) for each program stored on the server and these can be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, can indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

Likewise, the VOD video server 22 dynamically allocates PIDs for the shadow packets associated with the respective audio and video component streams for each session. This information is included in the PSI sent in sessions requested by non-legacy clients. Just like in the more general composite storage architecture discussed in the previous section, the video server manages multiple resources and PIDs. The hybrid topology reduces the unique entities by one from eight to seven: there is no need for alternative ECM PID or data resource in the stored composite stream. This information will be added later in a downstream device providing the alternative conditional access encryption for those sessions destined for decoding upon a non-legacy client.

RE-ENCRYPTED DISTRIBUTION

A hybrid approach is provided in a re-encrypted distribution architecture. This approach is described in U.S. Patent Application serial no. 10/764,202, filed January 23, 2004 to Pedlow Jr. et al., entitled "Re-Encrypted Delivery of Video On Demand Content", which is hereby incorporated by reference. This topology leverages the paradigms established for pre-encrypted content preparation, storage, management, etc. but adds support for session based encryption for the alternative conditional access systems added to an existing incumbent system. Referring to the exemplary embodiment of **FIGURE 5**, a legacy decryption device 182, operating to decrypt using the legacy CA system 184, is added to the transport stream path exiting the VOD video server 22 (via

routing matrix 186). After the decryption device 182, the transport stream passes through a contemporary session based encryption device 188. The VOD session manager 26, on a session-by-session basis, determines which sessions will pass through the decryption device 182 intact and be modulated and transmitted to the subscriber unaltered. A path 190 between the routing matrices preserves the pre-encrypted content and delivers it to subscribers having legacy equipment. In either case, the output stream passes through routing matrix 152 to the appropriate edge resources for delivery to the subscriber terminal 50.

Alternatively, the VOD system session manager 26, through interaction with both legacy CA system 184 and alternate CA system 194, can both actuate the decryption device 182 and activate session based encryption device 188 for a particular session, thereby supporting subscribers with non-legacy equipment at their premises. Thus, this system 180 can support either legacy or non-legacy (alternate CA) encryption.

Certain embodiments of this architecture support pre-encryption on legacy systems not presently supporting session-based encryption, while providing the ability to deliver session based encryption for the alternative CA system 194 integrated into the existing legacy network. Certain embodiments of this architecture may face the same issues as mentioned previously and common to other pre-encryption topologies. In addition, it experiences the additional cost burden of a legacy decryption element and the challenges of dynamically configuring and operating such a device. There may be additional costs faced in a specific deployment for switching and routing equipment that may be necessary to move transport streams "around" the legacy decryption device.

Changes can be made to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager 26 can be made aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information is in turn transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server can be made to be cognizant of the conditional access resources (ECMs) for each program stored on the server and these can be dynamically allocated on unique PIDs

along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

In this example, the same movie using 3.618GB of storage in the clear VOD state would require 3.618GBytes to store using re-encryption supporting two different CA systems.

DYNAMIC COMPOSITION PRE-ENCRYPTION

Another pre-encrypted VOD architecture is dynamic composition pre-encryption. In this scheme, each program or movie is stored in three or more elements on the VOD video server 22. Referring to **FIGURE 6**, clear content is stored at 200. Critical packets are selected according to a suitable selection criterion associated with the selective encryption process. Thus, the content that is stored has either "critical" packets or non-critical packets. The "critical" packets generally constitute approximately 2% to 10% of the program (depending upon program content and the selection criteria used to select packets for encryption) and are encrypted. A separate copy of the critical content is maintained for each conditional access system supported by the MSO. In this illustration, for example, the critical packets associated with a first CA system (CA1) is stored at 202 while encrypted content associated with CA 2 is stored at 206. By using a selection criterion that involves selection of certain I-frames, the fast forward I-Frames can be made to incorporate the encrypted content and stored together as encrypted I-frames 210 (and 206). The packets in both the "critical" packet fast forward file 210 as well as the clear (unencrypted), non-critical packet file 200 are indexed to maintain temporal correlation between the two files. These indices either may be monotonic packet counts from start of stream or calculated packet offsets from the last PCR.

When a subscriber session is initiated, the main file 200 containing the clear content, less "critical" and fast forward packets, is queued in the video server for playout. In addition, the file containing the "critical" and fast forward packets 210, pre-encrypted in the CA format appropriate for the CPE of the subscriber requesting the session, is also

queued for playout. When the program playback is started, the video server reconstructs a single program multiplex in its stream buffer feeding the outgoing transport the correct sequence of packets based upon the indices in the two component files. Although, in general, only about 2-10% of the packets are encrypted in a selective encryption system according to the above pending patent applications, even further security is provided by encryption of all of the I frames in the present embodiment. Rewind I-frames can be stored either as encrypted or unencrypted packets.

While the external composition and data flow appears similar to the clear VOD system depicted in **FIGURE 1**, the internal architecture of the video server changes significantly, as shown in the exemplary storage architecture of **FIGUREs 6-7**.

Certain embodiments of this method offer several distinct advantages that may not be readily apparent. The stream files containing "critical" packets may be the same one as the extracted subfile containing all I-frames for "trick" modes, as was described previously in the general discussion of VOD system architecture. If this opportunity is taken, then a storage economy can be realized over <u>all</u> pre-encrypted schemes including traditional (unencrypted) VOD, as deployed today. The traditional VOD video server has three files for each feature or movie: two containing just I-frames (one in reverse order) and one containing the complete original copy. Research on encoded streams conducted by Sony has shown that the I-frames typically represent between 12%-21% of the total content, typically around 17%. With the dynamic composition method, if the "critical" packet files are chosen to contain complete I-frames, a separate file of critical data used solely for encryption purposes is no longer necessary, saving 2% to 10% storage for this method. In addition, since this method removes the redundant I-frames from the clear stream file, an additional (nominal) 17% storage savings is also realized. This indicates a potential 27% nominal (31% maximum) video server disk storage savings for a single CA system model over the composite storage model VOD system described above.

When compared to the segregated storage model described above, one entire duplicate copy of a program can be eliminated and the addition of <u>one</u> additional CA format adds **<u>no</u>** storage or bandwidth overhead when compared to a traditional clear VOD server implementation. The reason for the "free" second CA format is that the 17%

nominal storage saving realized by using the same I-frame file for both fast forward "trick" modes and "critical" content used for selective encryption is consumed by replicating just the I-frame file and encrypting it with the alternative CA format.

DYNAMIC COMPOSITION PRE-ENCRYPTION WITH FORWARD AND REVERSE INDEXING

This concept is explained in greater detail in U.S. Patent Application Serial No. 10/764,011, filed January 23, 2004 to Pedlow et al. entitled "Bi-Directional Indices for Trick Mode Video On Demand", which is hereby incorporated by reference.

If one takes the concept of dynamic composition pre-encryption described above one step further, the current convention in VOD systems to store the same I-frames of a movie in forward and reversed sequence to allow fast forward and rewind "trick" modes can be eliminated. An illustration of this concept is shown in the example of **FIGURE 7**. These dual files for forward and reverse are replaced by a single file 220 of I-frames in normal forward sequence with two sets of indices, one set 222 for playing the I-frame file in forward order and one set 224 for playing the I-frame file in reverse order. The appropriate sets of indices are chosen depending on whether forward or reverse high-speed motion is desired. The forward indices are also used to reconstruct the normal speed stream when matching the I-frame file to the non-critical content file to reconstruct the entire stream. On a clear or re-encrypted VOD system, this will allow up to about 21% storage savings. On a composite pre-encrypted storage system, up to about 42% storage savings may be realized

If the "trick" mode subfile and the "critical" data encrypted content file can be the same, the content is selectively encrypted at approximately a nominal 17% level, much higher than the commonly proposed Passage™ encryption level of approximately 2%, but carrying no inherent storage or system capacity costs, as do other schemes. For this system to work, some changes to the video server software design might be necessary, but these changes would be modifications to the existing processes and would not require substantial new development on the part of the server vendor.

A preprocessor can be used to perform selective encryption of content to be loaded onto the VOD video server 22. A modified file protocol can be used to allow the VOD video server 22 to import and associate these files. Either the preprocessor or the VOD video server 22 can be used to perform the indexing. An alternate instantiation can be used to perform all selective encryption pre-processing within the video server itself. This can be accomplished by modifying the video server application to add a pre-processor task as a separate executable, called by the server during the process to prepare content for pre-encryption.

Additionally, in certain embodiments, this method overcomes the classic pre-encryption issue of supporting trick modes, but retains the other common problems of encryption "shelf life" and the additional handling required to prepare the stream for use on the VOD system.

Changes to the method employed by the VOD system for creating dynamic PSI data can be used to implement this architecture. The VOD system session manager 26 is made to be aware of which conditional access method is appropriate for a session requested by a specific subscriber in order to select the appropriate "critical" data file for the session. This information is in turn transferred to the VOD video server 22 that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The VOD video server 22 is cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, indicates the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

If one refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require 3.159GBytes to store using dynamic composition pre-encryption supporting two different CA systems – a savings of almost 0.5 GB.

SESSION-BASED ENCRYPTION VOD DISTRIBUTION

In session based encryption, a basic premise is that a classic (clear) VOD server 22 such as shown in **FIGURE 1**, is modified to add an encryption device in series with the transport stream between the video server 22 and the QAM modulator of 46. In certain embodiments, the encryption device may be integrated with the QAM modulator 46 and/or other components. The commercially available Scientific-Atlanta MQAM and Harmonic NSG products are commercial examples of such devices.

The outgoing transport stream, containing multiple, independent VOD sessions and serving multiple subscribers, is encrypted at the point of distribution to the plant and in turn to the subscribers. The control of the encryption and entitlements is based upon interaction between the session manager 26, which controls the session, video server 22 and the conditional access system through defined interfaces. Many session based VOD architectures share the following common drawbacks:

- Coordination and/or distribution of entitlements and synchronization between session manager, conditional access system and stream encryption device.
- Security of the clear content from theft or piracy before loading on the video server and while stored in the system.
- Additional costs for adding both legacy and alternate stream encryption devices.
- Availability of legacy stream encryption devices with reasonable densities (session capacity).
- According to MSOs familiar with the subject, session based VOD streams are unsupported by certain existing conditional access technologies.
- With session-based encryption (compared to the pre-encryption scheme) additional security is afforded by the application of unique encryption keys used for every session of the same program.

In most cases, the video server does not need to generate special PSI that is aware of the conditional access method used for a specific session. The encryption device(s) downstream of the video server will append CA information specific to each session processed at the time/point of encryption. The VOD session manager 26 manages which

streams are processed by which CA method and in some cases, manages dynamically routing the streams to/through the encryption devices appropriate for a particular session.

As with other architectures, there are variations on the basic architecture of the session-based system and some of those variations are described below.

SEGREGATED SESSION BASED ENCRYPTION

Segregated session encryption is the extension of session-based encryption to multiple conditional access systems operating in conjunction with a single VOD system. An exemplary architecture of a segregated session based encryption system 240 is depicted in **FIGURE 8**. System 240 includes provisions for providing the appropriately encrypted stream for a specific subscriber session by routing the outgoing stream from the VOD video server 22 to the subscriber terminal 50 on a transport stream and resultant RF carrier, carrying only a single common conditional access format. Sessions using other conditional access formats are similarly constrained (segregated) to other homogeneously encrypted transports/carriers. There is no sharing of resources between the CA systems and they operate independently.

During initiation of a new session, the VOD session manager 26 determines which conditional access format is used by the requesting subscriber terminal 50 based upon information received either directly from the subscriber terminal 50 or from another source, such as the billing system or other database. The VOD session manager 26 then determines the path to the appropriate encryption resource(s) 184, 188 and 194 having access to an RF node serving the subscriber's service area. This is done in a similar manner to the method used in large centralized clear VOD systems to find the appropriate video server(s) that can deliver a stream to the requesting subscriber. Once an appropriate route is determined, routing matrix and re-multiplexer 242 and routing matrix 244 respectively provide the appropriate routing.

Once a solution to the routing matrix is determined, the session manager 26 coordinates the configuration of the routing elements and directs the CA system to apply encryption to the session through references to the assigned transport resources (PIDs).

This system presents a complex, real-time management requirement for determining usable resources available to apply to a new session and available spectrum transport slots. It uses equipment to perform stream routing (switch fabric) between the VOD video server(s) 22 and the encryption devices 188, though these capabilities might be available integrated into other elements of the system. Additional spectrum is used to maintain segregation of the sessions on homogeneously encrypted transport streams and carriers.

A segregated session based encryption scheme such as 240 uses, to some varying degree, duplication of encryption resources such as encryption devices 188, since support of simultaneous sessions in differing conditional access formats is required. Careful traffic modeling can optimize the tradeoff between system capacity/resource availability and capital expenditure.

If one again refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require 3.618GBytes to store using segregated session based encryption supporting two different CA systems. The system can be optimized in a manner similar to that described in the section describing dynamic composition based pre-encryption. One I-frame file can be removed for rewind and a dual set of indices created for the remaining I-frame file to support both forward and reversed video sequences. In doing so, the total storage required for the example movie could be reduced to 3.159GBytes.

COMPOSITE SESSION BASED ENCRYPTION

The composite session based encryption approach (another session based approach) is similar to the segregated approach except that the transport streams/carriers provided to subscribers are heterogeneously encrypted and is depicted embodied in system 260 of **FIGURE 9**. A single transport may contain any combination of two or more conditional access formats operating independently on an MPEG program basis, representing individual subscriber sessions.

This scheme eliminates some of the complex real-time resource management processes used to determine available encryption resources, but instead trades it for the requirement that encryption resources appear in matched sets.

The VOD session manager 26 determines which CA format is appropriate for a given subscriber session and determines a VOD server 22 that has access to the node representing the subscriber's service area. It then activates the appropriate CA resource in the encryption "set" attached to the node. It is noted that a process such as the Passage™ process of selective encryption is not employed, since there is never an opportunity to share any common content between subscriber sessions in a VOD paradigm. A technical consideration that should be considered is the configuration of systems with specific combinations of legacy encryption and/or remultiplexing equipment. This is especially true if the alternative encryption is embodied within the device performing the remultiplexing. The Harmonic NSG is a commercial product that can be used for this purpose. If the legacy system transmits data on unannounced PIDs or has critical latency concerns, this may be problematic if the device performing re-multiplexing is not aware of these requirements.

FIGURE 10 depicts the encrypted content flow from stream files to composite stream in a composite session based encryption system. Stream files 264 are processed by legacy encryption device 266, while stream files 270 are processed by the alternate CA encryption device 274. The output streams from encryption devices 266 and 274 are multiplexed at stream re-multiplexer 280 to produce the composite stream as an output to the subscribers.

If one again refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require 3.618GBytes to store using composite session based encryption supporting two different CA systems. The system could be optimized in a manner similar to that described in the section describing dynamic composition based pre-encryption. One I-frame file would be removed for rewind and a dual set of indices created for the remaining I-frame file to support both forward and reversed video sequences. In doing so, the total storage required for the example movie could be reduced to 3.159GBytes.

BATCH-BASED ENCRYPTION VOD DISTRIBUTION

The concept of batch-based encryption for VOD distribution, another session based mechanism, as depicted in **FIGURE 11** represents many of the best aspects of both session and pre-encrypted architectures. As can be seen in **FIGURE 11**, the batch based VOD system has a topology different from the other systems presented in this document.

The content can be stored entirely in the clear on the VOD video server 22, similar to the session-based system, but is contained in two files, representing "critical" packets and non-critical packets, just as in the case of the dynamic composition architecture. Likewise, the same opportunities for storage efficiency are available if the "critical" packet files are also used as the "trick" mode I-frame files 314 and 316 as shown. However, unlike the dynamic composition architecture, the "critical" packets are stored in RAID files 70 unencrypted. Additionally, this scheme departs from the dynamic composition architecture because there is no requirement to maintain an independent copy of the "critical" packet file for each conditional access system supported, providing further, substantial storage savings over the other architectures, typically on the order of approximately 12% to 21% per conditional access system supported if the critical packets are the same as the I Frames. Otherwise, the critical content may be approximately 2 – 10 percent.

The Fast Forward I frames are generally not encrypted. So, critical content is identified and stored in a separate file. Once the encryption format is determined, data are burst in at a maximum rate through the appropriate encryption device and stored in temporary storage 310. Regular non-critical content is sent to 324. From there on, critical and non-critical content is spooled and multiplexed in by the Multiplexer and Buffer 84 for sending to the STB at the regular transport rate.

Thus, in this example, a Video On Demand (VOD) method consistent with certain embodiments involves processing content to be delivered in a VOD method by selecting first portions of the content for encryption under a selective encryption system and selecting second portions of the content to remain unencrypted. The first and second portions are stored until receiving a request for delivery of the content, the request being

from a terminal having decryption capabilities associated with a first decryption method. The first portions are then bulk encrypted in bulk legacy encrypter 302 to produce encrypted first portions stored at 320. The encrypted first portions are stored in a temporary store 310. The second portions are queued to temporary storage at 324 for delivery to the terminal. A stream of selectively encrypted content is assembled from the encrypted first portions and the second portions.

FIGURE 12 shows a system similar to that of **FIGURE 11** except that this system utilizes the dual indexing to provide trick play as described in conjunction with **FIGURE 7.**

When the session manager 26 initiates a new session at the request of a subscriber terminal 50, the encryption technology appropriate for the subscriber's equipment is determined. The file for the selected feature containing the clear, non-critical content is queued in the VOD video server 22 for playout. In addition, a second file, containing the clear stream of "critical" packets is accessed; its contents are immediately streamed through a dedicated port on the VOD video server 22 at the maximum sustainable transport medium data rate (1Gbit/S for Gig-E, 200+Mbit/S for ASI, 38.8Mbit/S for DHEI) directly to the encryption resource identified by the session manager, i.e., either legacy encryption device 302 or alternate CA encryption device 306. This burst transferred file of I-frames, constituting only about 12% to 21% of the video frames (assuming the critical content equates to the I Frames) in the program is bulk encrypted at the highest rate that the encryption device 302 or 306 and transport media can sustain. The encrypted I-frame content that emerges from the encryption device is captured to either a RAM or disk buffer 310 resources within the VOD video server 22. For a 2-hour movie, with a nominal 17% (approximately) I-frame content, this would require 450Mbytes of temporary storage per session.

When the program playback is started, a multiplexer reconstructs in the stream buffer feeding the outgoing transport the correct sequence of packets based upon the indices in the clear, non-critical content component file and the smaller, batch-encrypted content that was captured back to the VOD video server, as described above.

This architecture, in addition to the storage efficiencies described both under the dynamic composition architecture description as well as in the previous paragraph, offers additional, significant advantages in certain embodiments. The batch encryption of "critical" packet files can allow for a significant reduction in the number of encryption devices required to provide encrypted delivery of VOD content. If one assumes support of two independent conditional access systems using this architecture, the I-frame and critical data residing in the same file and using a typical I-frame overhead (~17%), then a single pair of encryption devices (incumbent & alternative) can support the same number of sessions as 60 pairs of encryption devices in a session based architecture (60:1). Another advantage over the pre-encryption scheme, shared with session-based encryption, is the additional security afforded by the application of unique encryption keys used for every session of the same program.

An alternate embodiment pre-encrypts sessions of I-frames and stores them in the buffer for later consumption. In this manner, there would be no latency to delivering a new session due to the time overhead required to batch encrypt the file. The buffer of pre-encrypted I-frames could be replenished in the background to maintain a constant "inventory" of available sessions for delivery.

The method employed by the VOD system for creating dynamic PSI data can be modified to implement this architecture. The VOD system session manager 26 can be made aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information is in turn transferred to the VOD video server 22 that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The VOD video server 22 can be cognizant of the conditional access resources (ECMs) for each program stored on the server and these are dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, must indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

If one again refers to the example movie scenario described above, the same movie using 3.618GB of storage in the clear VOD state would require 2.700GBytes to store using batch-based encryption supporting two different CA systems.

Thus, a Video On Demand server arrangement, consistent with certain embodiments, receives content from a selective encryption processor that processes content to be delivered in a VOD method by selecting first portions of the content for encryption under a selective encryption system and selecting second portions of the content to remain unencrypted. At least one computer readable storage device is provided. The processor stores the first and second portions in the at least one computer readable storage device. When the processor receives a request for delivery of the content, the request being from a terminal having decryption capabilities associated with a first decryption method, it is programmed to send the first portions to an encrypter that encrypts the first portions using a bulk encryption process to produce encrypted first portions. The processor then stores the encrypted first portions in a buffer and queues the second portions for delivery to the terminal. The processor is programmed to then assemble a stream of selectively encrypted content from the encrypted first portions and the second portions.

Referring now to **FIGURE 13**, a flow chart depicts a process 400 for batch based encryption of VOD content consistent with certain embodiments starting at 404. In this embodiment, the content (movie, etc.) is stored in the VOD server 22 on two separate files at 408– one for content designated as clear, and one designated for encryption. In a selective encryption process, certain selected portion of the content is encrypted while the remaining content remains unencrypted. The first file contains all content that is designated in a selective encryption selection process to remain unencrypted. The second file contains all content that is designated in the selective encryption selection process to be encrypted. At 412, when a request for the content is retrieved, the content designated to remain clear is spooled up from the first file to a queue at 416 for delivery to the subscriber.

In the present example, two possibilities for encryption are available, namely CA1 and CA2 (which, for example, can represent legacy encryption and an alternative

encryption process). The subscriber's STB is likely only capable of decryption under CA1 or CA2, so a determination is made at 420 as to which conditional access system is associated with the STB making the request for content. This can be accomplished in any number of ways, including providing that information at the time of a request for content or looking up that information in a database (e.g., a database stored at the billing system 34 or elsewhere). If the request comes from a STB that is enabled for decryption of CA1 encrypted content, the second file containing the content designated by the selective encryption algorithm for encryption is bulk encrypted at 424 using the CA1 encryption system. If, on the other hand, the request comes from a STB that is enabled for decryption of CA2 encrypted content, the second file containing the content designated by the selective encryption algorithm for encryption is bulk encrypted at 432 using the CA2 encryption system.

In either case, the encrypted content from 424 or 432 is passed to 436 where the encrypted content is buffered up for playout to the requesting STB at 436. At 440, the content is played out by streaming the content to the requesting STB. This is accomplished by reconstructing a complete content stream from the clear content queued up from the first file and the encrypted content from the second file. The content can then be streamed to the recipient STB at 440 and the process ends at 448.

Thus, a Video On Demand (VOD) method consistent with certain embodiments involves processing content to be delivered in a VOD method by selecting first portions of the content for encryption under a selective encryption system and selecting second portions of the content to remain unencrypted; storing the first portions; storing second portions file; receiving a request for delivery of the content, the request being from a terminal having decryption capabilities associated with a first decryption method; encrypting the first portions using a bulk encryption process to produce encrypted first portions; storing the encrypted first portions in a buffer; queuing the second portions for delivery to the terminal; and assembling a stream of selectively encrypted content from the encrypted first portions and the second portions.

A Video On Demand (VOD) method, consistent with certain embodiments involves processing content to be delivered in a VOD method by selecting first portions

of the content for encryption under a selective encryption system and selecting second portions of the content to remain unencrypted; receiving a request for delivery of the content, the request being from a terminal having decryption capabilities associated with a specified decryption method; encrypting the first portions under a specified encryption method using a bulk encryption process to produce encrypted first portions, the specified encryption method being selected from one of a plurality of encryption methods; storing the encrypted first portions in a first file; storing second portions in a second file; storing the encrypted first portions in a buffer; queuing the second portions for delivery to the terminal; and assembling a stream of selectively encrypted content from the encrypted first portions and the second portions. The first and second portions may be stored in a VOD server in first and second files in accordance with certain embodiments.

Those skilled in the art will recognize, upon consideration of the above teachings, that certain of the above exemplary embodiments are based upon use of a programmed processor serving, for example, as video server or servers 22 or session manager 26. However, the invention is not limited to such exemplary embodiments, since other embodiments could be implemented using hardware component equivalents such as special purpose hardware and/or dedicated processors. Similarly, general purpose computers, microprocessor based computers, micro-controllers, optical computers, analog computers, dedicated processors, application specific circuits and/or dedicated hard wired logic may be used to construct alternative equivalent embodiments.

Those skilled in the art will appreciate, upon consideration of the above teachings, that the program operations and processes and associated data used to implement certain of the embodiments described above can be implemented using disc storage as well as other forms of storage such as for example Read Only Memory (ROM) devices, Random Access Memory (RAM) devices, network memory devices, optical storage elements, magnetic storage elements, magneto-optical storage elements, flash memory, core memory and/or other equivalent volatile and non-volatile storage technologies without departing from certain embodiments of the present invention. Such alternative storage devices should be considered equivalents.

Certain embodiments described herein, are or may be implemented using a programmed processor executing programming instructions that are broadly described above in flow chart form that can be stored on any suitable electronic or computer readable storage medium and / or can be transmitted over any suitable electronic communication medium. However, those skilled in the art will appreciate, upon consideration of the present teaching, that the processes described above can be implemented in any number of variations and in many suitable programming languages without departing from embodiments of the present invention. For example, the order of certain operations carried out can often be varied, additional operations can be added or operations can be deleted without departing from certain embodiments of the invention. Error trapping can be added and/or enhanced and variations can be made in user interface and information presentation without departing from certain embodiments of the present invention. Such variations are contemplated and considered equivalent.

While certain embodiments herein were described in conjunction with specific circuitry that carries out the functions described, other embodiments are contemplated in which the circuit functions are carried out using equivalent software or firmware embodiments executed on one or more programmed processors. General purpose computers, microprocessor based computers, micro-controllers, optical computers, analog computers, dedicated processors, application specific circuits and/or dedicated hard wired logic and analog circuitry may be used to construct alternative equivalent embodiments. Other embodiments could be implemented using hardware component equivalents such as special purpose hardware and/or dedicated processors.

Software and/or firmware embodiments may be implemented using a programmed processor executing programming instructions that in certain instances are broadly described above in flow chart form that can be stored on any suitable electronic or computer readable storage medium (such as, for example, disc storage, Read Only Memory (ROM) devices, Random Access Memory (RAM) devices, network memory devices, optical storage elements, magnetic storage elements, magneto-optical storage elements, flash memory, core memory and/or other equivalent volatile and non-volatile storage technologies) and / or can be transmitted over any suitable electronic

communication medium. However, those skilled in the art will appreciate, upon consideration of the present teaching, that the processes described above can be implemented in any number of variations and in many suitable programming languages without departing from embodiments of the present invention. For example, the order of certain operations carried out can often be varied, additional operations can be added or operations can be deleted without departing from certain embodiments of the invention. Error trapping can be added and/or enhanced and variations can be made in user interface and information presentation without departing from certain embodiments of the present invention. Such variations are contemplated and considered equivalent.

While certain illustrative embodiments have been described, it is evident that many alternatives, modifications, permutations and variations will become apparent to those skilled in the art in light of the foregoing description.

What is claimed is: